

CHAPTER 3

IP NETWORK TUTORIAL

3-1 BASEWIDE INTERNET PROTOCOL (IP) NETWORK

The basewide UMCS network will be based on Ethernet and Internet Protocol (IP) standards. IP and Ethernet are more than adequate for the data transmission requirements of a UMCS. In addition, they are widely-adopted, mature technologies readily available from many vendors and supported by a huge number of contractors, as well as on-site personnel at most installations. IP networks are peer-to-peer networks where all nodes (UMCS servers, UMCS workstations, and Building Point of Connection (BPOC) hardware) have the same level of control over the communications and can control their own activities.

The most important consideration when developing a basewide UMCS network is coordination with the installation Directorate Of Information Management (DOIM). This can be beneficial in several ways:

- They may allow direct access to the basewide Information Technology (IT) network. A properly designed UMCS system does not require a great deal of bandwidth (see below for further discussion of this topic). This will most likely result in the UMCS being placed on a separate VLAN (Virtual LAN) on the basewide network.
- Many installations have spare physical media (typically fiber-optic cable) installed and the UMCS network may be able to use this media. In this case, the UMCS network will be a logically separate network (from the basewide IT network), but will share physical resources co-located with the basewide IT network.
- Even if the UMCS contractor is required to install completely separate networking hardware, the DOIM can be a valuable source of information regarding local support for networking technologies.

3-2 BACKGROUND

3-2.1 OSI Reference Model

The many aspects of a network are frequently represented as layers as laid out in the ISO Standard 7498. (ISO 7498-1984 Information processing systems- Open Systems Interconnection (OSI)- Basic Reference Model) That document presents a 7-layer model of information exchange as shown in table 7-1. Most descriptions of computer networks, especially open networks, are based on this reference model. The layers can be thought of as steps in the conversion of a message from something with meaning at the application layer, to something measurable at the physical layer (i.e voltages), and back to meaningful information.

The layered approach to network design is valuable because it allows developers to take advantage of existing standards (such as IP at the network layer or Ethernet at the physical layer) without becoming tied to one technology.

Table 3-1. OSI Reference Model

| Layer Number and Name | | Layer Description |
|-----------------------|--------------------|-----------------------------------------------------------------------------------------------------------------|
| 7 | Application Layer | Window between applications and the network. |
| 6 | Presentation Layer | Coordinates representation of information between different applications. |
| 5 | Session Layer | Synchronizes and structures data exchange between specific users. |
| 4 | Transport Layer | Responsible for end-to-end reliability of data transfer. |
| 3 | Network Layer | Addressing and routing independent of media and topology. |
| 2 | Data Link Layer | Responsible for point-to-point reliability. Media access. Representation of bits and bytes as physical signals. |
| 1 | Physical Layer | Electrical characteristics of devices and conductors. |

3-2.2 Copper Wires, Fiber Optics, And Other Physical Media

A variety of media are used at the physical layer. Coaxial cable is used in older installations, however, it's use is not recommended for new installations.

For some lower-speed networks, the choice of wire type is not critical. Some commonly used types include Unshielded Twisted Pair (UTP), Shielded Twisted Pair (STP) and even voice-grade telephone wire. For higher speed networks (over 10 MBPS), the choice of wire becomes more critical. Wire meeting EIA/TIA 568 Category 5 (most commonly called "Cat 5 cable") is commonly used for this purpose.

3-2.2.1 Fiber Optics

Fiber optic cable consists of small fiber cores encased in a thin jacket referred to as the cladding, which is in turn enclosed by a thicker plastic or Teflon jacket. A light source (a laser diode or light-emitting-diode) at one end of the cable introduces coded light pulses into the fiber. The light pulses are transmitted through the fiber to a photo diode at the other end, which receives the light pulses and converts them to electrical signals. Fiber optic cables (but not fiber optic equipment) can be installed in explosive and flammable environments. The bandwidth of this media is virtually unlimited, and extremely high data transmission rates can be obtained. The signal attenuation of high quality fiber optic cable is very low. Fiber optics is a reliable communications media best suited for high speed inter-building data transmission. Fiber optics is immune to radio frequency electromagnetic interference, and does not produce electro-magnetic radiation emission; hence, fiber optics can be used in secure areas. Fiber optic terminations are more expensive than copper and fiber is seldom cost-effective for short distances.

Unlike with wires, optical fibers terminate at a light source or a light receiver, never both. Because of this, light only travels in one direction down a fiber and bi-directional communication requires a pair of fibers. Fiber optic cables come in a variety of sizes

and terminations. The three most common sizes are 62.5/125 micron MultiMode Fiber (MMF), 50/125 micron MMF, and 10/125 micron SingleMode Fiber (SMF). The two numbers given for the size refer to the diameter of the core and cladding, respectively. Single Mode Fibers have a higher bandwidth than Multi Mode Fibers, particularly for long distance runs, but their smaller size results in a higher installed cost due to more expensive terminations. For UMCS systems, the difference in bandwidth is irrelevant; either technology will have sufficient bandwidth for all but the longest (5 km or more) runs. There are a wide variety of incompatible connectors for performing fiber terminations, including ST, SC, FDDI, MTJR, and FC. ST was the most commonly used type and is still found in many existing fiber installations, along with SC. New construction favors MTJR or duplex SC connectors (that is, connectors where two connectors are physically joined as one unit).

The use of fiber optics equipment and connectors will introduce optical signal losses/gains that must be accounted for during the design. The contractor will calculate the optical flux budget/gain during UMCS data transmission system design.

3-2.2.2 Wireless

Two common methods of wireless networking are Radio Frequency (RF) and Free Space Optical (FSO).

3-2.2.2.1 Radio Frequency (RF)

RF communication is inexpensive and easy to use. It can be used with mobile equipment. It's main limitations are short range and the potential for interference. Any use of RF on and installation should be carefully coordinated with the DOIM, communications officer, or other responsible authority.

3-2.2.2.2 Free Space Optical (FSO)

FSO is a fixed point-to-point link between two optical transceivers. The transceivers are mounted on rooftops or free-standing towers and pointed at each other. FSO communication is easy to use and capable of high bandwidths. Because it is a point-to-point link, it is immune to outside interference and it is generally fairly easy to get approval from the installation to use. At longer ranges and higher bandwidths, it can be affected by weather conditions, particularly fog.

3-2.3 Network Interface Cards And Transceivers

To function on a network, end equipment (DDC hardware, front-end PCs) must have a network interface of some sort, which will handle functions at layers one and two (physical layer and data link layer). For a PC, this will be a network interface card (NIC), either plugged into a slot in the computer or integrated onto the PC motherboard.

A transceiver is a physical layer device that generates the proper voltage levels for the cable/protocol in use. Transceivers have a receiver, which listens for signals, and a transmitter which can transmit a "1" or a "0" or not transmit. Transceivers are typically part of a hardware interface that includes additional circuitry for generating the proper bit rate and additional bits needed by the protocol. Transceivers are the part directly connected to the wire.

3-2.4 Bus Arbitration

Bus arbitration addresses the issue of having multiple transmitters coexist on a single wire. A common issue in network design is bus arbitration -- how to prevent multiple transceivers from transmitting at the same time on the same network and generating what is known as a "collision". Three common approaches are token-passing, polling, and CSMA/CD (Carrier Sense Multiple Access / Collision Detection). Polling is generally not desired and Ethernet uses CSMA/CD.

3-2.4.1 Token-Passing Network

In a token-passing network, a special message is passed from transceiver to transceiver, which says, in effect "you have the floor" (think of a microphone passed around at a meeting). Transceivers can transmit only if they have the token. Different implementations determine how long a transceiver can have the token, how many messages can be transmitted before the token must be given up, or even which transceivers are eligible to have the token. Some variations allow a transceiver without the token to transmit in response to a direct query from a transceiver with the token (i.e. transceiver "A" has the floor, but can ask a question of, and get a response from transceiver "B" without giving up the floor). Master-Slave/Token-Passing, MS/TP, a common network protocol, specifies two types of transceivers, Masters who can get the token, and Slaves, which cannot and can only respond to transmission requests from a Master transceiver. (In this case, Masters can talk to Slaves and each other; Slaves cannot talk to other Slaves.) The intent of this division was to promote interoperability by making proprietary devices Slaves and interoperable devices Masters, thus avoiding the problem of having two incompatible Slaves talking to each other, or even one incompatible Slave broadcasting "gibberish" (i.e. an incompatible protocol) on the network.

3-2.4.2 Polling Network

A polling network can be thought of as a token network where the token always remains with one transceiver. In effect, there is a single Master, who can (and must) talk to every device on the network. Other devices can only answer requests from the Master device. This is called a polling network because in effect the Master has to "poll" each device under it's control to obtain data from that device. Polling networks have performance limitations, particularly with respect to bottlenecks at the Master device and failure of the entire network if the Master device fails.

3-2.4.3 CSMA/CD Network

If a polling network is an orderly meeting, a typical Carrier Sense Multiple Access / Collision Detection (CSMA/CD) network is more like a typical "bull" session, with everyone free to talk at once. In a CSMA/CD network, each transceiver is designed to follow 2 rules:

- "Do not transmit when there is a signal on the network already (i.e. another transceiver is transmitting). This is the Carrier Sense Multiple Access rule: Anyone can transmit, but only when the line is quiet."
- "If you are transmitting and you detect a signal from another transmitter, cease transmission, wait a short (and random) time, then re-transmit. This rule is necessary for the case where two transceivers start transmission at the same time. This is the Collision Detection rule: In case of a collision, stop and try again."

This is the most normal conversations work: You only talk when no one else is talking; if two people start talking at once, they both shut up, then one of them starts up again. A drawback of CSMA/CD networks is that they seldom achieve anywhere near their theoretical bit rates. As more devices try and push more data through the bus, more collisions are generated, which results in more re-transmissions, which results in more data on the bus, which causes more collisions, etc. (picture a meeting with everyone talking at once - a CSMA/CD network has no recovery mechanism for someone to take charge and yell "Shut up"). Ethernet (the most obvious example of a CSMA/CD network) usually saturates at about 30% - 40% of its theoretical bandwidth (i.e. a 10 MBPS Ethernet can carry about 3 - 4 MBPS of data) under normal conditions, with multiple devices trying to transmit data at once (it can approach its theoretical limit when there is a single transmitter and a single receiver).

3-2.5 Segment Length Limits

A segment is a single cable as part of a network. It's the smallest "unit" of a network. All networks have limitations on how long a single segment can be. In some cases, these limits are related to signal loss in the cable, for other networks the physical limits are due to timing issues related to collision resolution. This is due to the finite propagation velocity of a signal down the wire and the timing issues associated with rule #2 above. (If you have ever had a long distance phone conversation where one party is using Voice-Over-IP (VoIP), you know that the delays in the process make it much more likely that the two parties will interrupt each other).

3-2.6 Collision Domains

A collision domain is one or more segments where a collision anywhere in the domain is a collision everywhere in the domain (i.e. any two transceivers on the domain will produce a collision if they transmit at the same time).

3-2.7 Addressing And Packets

With multiple devices comes the need for addressing. This is one of the functions of the Data Link Layer (Layer 2). There is also addressing at the Network layer (layer 3), which will be discussed later. Every Ethernet adapter has a MAC address - that's its Ethernet address. Every Lon network adapter has a NeuronID number - that's its low-level Lon address. Some devices permit changing of these Layer 2 addresses, but most often they are fixed in hardware.

Addresses are generally sent over the same media as the data. In addition to the data and the recipient's address, there are a number of other things that are typically sent at the same time; this collection of items is called a data packet. A packet will have:

- Destination address: (required, but may be the special address "ALL")
- Source address: (optional)
- Length of data: (optional) Some protocols allow for variable amounts of data in one packet - if so, the receiver needs to know how long this packet is.
- Checksum: (optional) More error correction
- Data: The data to transmit.
- Other protocol-dependent information.

Keep in mind that each of these pieces consists of individual bytes, which ultimately get sent as a series of bits which in turn are represented as voltages (of well-defined duration) on the wire. (or light pulses, in the case of fiber optics).

3-2.8 Repeaters, Hubs, Switches, Bridges, And Media Converters.

3-2.8.1 Repeater

A repeater is a device that performs signal regeneration (takes a signal on its input, cleans it up, and puts the signal out on its output). So a repeater connects multiple segments. Repeaters allow for longer cable runs if the limitation is due to signal degradation, but do not remove limits related to collision timing (i.e. segments connected by repeaters are in the same collision domain).

3-2.8.2 Hub

A hub is a multi-port repeater. It has multiple input ports and multiple output ports. A signal on any input port is sent to all output ports.

3-2.8.3 Media Converter

A Media Converter is a repeater that changes media types, i.e. from Copper to Fiber Optic. Because it simply passes bits from one side to the other, it cannot be used to connect media operating at different speeds. Media converters are commonly found in fiber networks to connect a hub or switch to a pair of optical fibers.

3-2.8.4 Switch

A Switch is much like a hub, but does not send signals from an input port to all output ports. Instead, a switch examines the destination address (for Ethernet, this is the destination MAC address) of the packet and only forwards it to the correct port to reach the destination. Switches are usually “learning switches”, which means they have the ability to query the network and determine the addresses of devices connected to them. Note that this addressing is at the physical layer (layer 2), if the address that the switch looks at is a Networking layer (layer 3) address, the device is called a router instead. Switches with only two ports are often called bridges.

Switches have 3 advantages over hubs:

- **Bandwidth:** Because the switch only copies data from one input to one output, another device can simultaneously transmit data from another unused input to another unused output – the switch keeps the two transmissions separate and there is no collision.
- **Security:** Because the switch only copies data from one input to one output, a malicious user ‘listening’ (monitoring network traffic) on another port (not involved in the traffic) cannot intercept the data
- **Cable Length Restrictions:** Most switches utilize a “store-and-forward” scheme. This means that the switch reads in the whole data packet, determines where to send it, then forwards the packet. The transmission of the packet by the switch is a separate transmission and collisions on the output port do not affect the input port. Because of this, ports connected to a switch are on separate collision

domains, and two or more segments connected by switches can exceed the length requirement of a single segment.

3-2.8.5 VLAN

Most modern switches are remotely managed and have the capability of supporting multiple VLANS (Virtual LAN). A managed switch is one that can have its switching capabilities configured; a remotely managed switch can be configured remotely, typically over the IP network. A normal switch is capable of forwarding an incoming packet to any output port (based on destination address), a switch supporting a VLAN can be configured to only allow communications between subsets of ports. For example, a typical 24 port switch with VLAN capability might be set up as follows:

| Name | Ports | Description |
|-------|-----------|------------------------------------|
| UMCS | 1,2, 8-15 | Basewide UMCS LAN |
| HR | 4-7 | Basewide Human Resources LAN |
| 3rdBn | 16-20 | 3 rd Battalion Training |
| DOIM | 24 | DOIM Control LAN |
| --- | 3, 21-23 | Unused |

In this example, there are 4 independent LANs sharing this switch; a malicious user on the HR LAN cannot access any network assets not connected to ports 4-7. Traffic on these LANs is kept separate by the configuration of the switch (recall that a normal switch will allow traffic between any two ports). The DOIM LAN is a special LAN set up specifically to control the switch itself; the switch is configured only to accept control commands from port 24. This configuration can be accomplished dynamically; the DOIM can easily move a port from one VLAN to another.

VLAN technology is a very powerful and very common means of controlling network bandwidth and access without requiring completely separate hardware. In particular, applications with greater security risks (perhaps machines in unsecure areas such as barracks) or with greater security requirements (not to include classified networks!) can safely co-exist on the basewide network. The DOIM can utilize a VLAN to allow the UMCS to reside on the basewide IT network while remaining segregated from the other networks on the base.

3-2.8.6 Bridge

A Bridge is similar to a media converter, in that it changes media types, but a bridge can also connect media operating at different speeds. Since bridges often incorporate a 'store-and-forward' strategy, the two segments connected via the bridge are in different collision domains. Also, many bridges perform filtering by physical address; in this sense they incorporate the functionality of a switch as well.

3-2.8.7 Terminology

These terms (repeater, hub, switch, bridge and media converter) are frequently misapplied. In particular, people often call switches "hubs" and bridges "media converters". Also, a single piece of hardware frequently incorporates multiple functions. For example, one may purchase a 24 port 10/100 Base-T switch with Gigabit Fiber uplink. This device will have 24 copper ports which will auto sense between 10Base-T (10 MBPS) and 100Base-T (100 MBPS) (depending on what speed equipment is connected to each port) and one fiber port at 1000 MBPS. It provides packet filtering by

physical address (as a switch) and network media and speed conversion between 10 MBPS, 100 MBPS, and 1000 MBPS (as a bridge).

3-2.9 Network Topologies

Network topologies fall into one of three classes:

- Bus topology is a configuration where the devices connect directly to the same media by means of connectors in a daisy chain configuration. As a special case of the bus, a ring topology is often used where the two ends of the bus are joined together. While a few obsolete networking technologies used this ring to pass data in a circle around the ring, modern networks use the ring to create a redundant path in case of a network break. Care needs to be exercised in setting up a ring to ensure that there is a 'preferred' path between two points on the ring.
- Star topology is a configuration in which interconnection hardware connects radially to multiple field equipment panels.
- Hybrid topologies which are hierarchical combinations of star and bus topologies.

To further complicate the issue, networks which at one level are actually star networks are often referred to as a bus network. This is because it is often easier to logically consider the network as a bus, even though the actual hardware implementation may be a star configuration. (See Figure 3-2)

3-2.10 Ethernet

An Ethernet network consists of physical media (copper wires and fiber optic (FO) cables), interface hardware (primarily network interface cards (NIC) and media converters (MC)), and interconnection hardware (primarily switches and hubs). All modern Ethernet hardware (except 10Base-2) utilizes point-to-point links, with interconnection hardware (generally hubs or switches, occasionally bridges or repeaters) at one or both ends of the link, so that the underlying topology is a star.

There are several different types of Ethernet, distinguished by media type and data transmission speed:

- 10Base-2. This network functions at 10 MBPS and is based on 50 Ohm coaxial cable connected in a bus topology or star topology. This is obsolete, seldom encountered, and should not be used for new construction
- 10Base-T. This network functions at 10 MBPS and is based on Unshielded Twisted Pair (UTP) wiring. This network can even use (with some limitations) regular phone lines for the physical media. Maximum segment length is 100 meters.
- 100Base-TX. This network functions at 100 MBPS and is based on UTP wiring which must conform to EIA/TIA 568 Category 5 or better, more commonly referred to as "Cat 5" cable. Maximum segment length is 100 meters.
- 10Base-FL. This is a 10 MBPS network utilizing fiber optic cables with a maximum segment length of 2000 meters.

- 100Base-FX. This is a 100 MBPS network utilizing fiber optic cables with a maximum segment length of 400 meters.
- 1000Base-T. This is a 1000 MBPS network utilizing Cat-5e wiring with a maximum segment length of 100 meters. While Cat-5 wiring can be used for 1000Base-T networks, existing wiring installations should be re-certified for 1000Base-T as some of the performance criteria are more stringent than for 100Base-TX. This, and 1000Base-SX / 1000Base-LX are often referred to as "Gigabit Ethernet".
- 1000Base-SX and 1000Base-LX. These are 1000 MBPS networks utilizing fiber with a maximum segment length of 220 to 5000 meters, depending on the exact implementation.

There is currently under development a standard for 10-Gigabit Ethernet.

Ethernet networks generally have hybrid topologies, with star networks in buildings and a bus network acting as the Ethernet 'backbone', joining all the buildings on the installation together.

Figure 3-1 shows a fiber backbone connecting several buildings together. Inside each building is a fiber switch / media converter to connect the fiber backbone to a 1000Base-T copper Ethernet segment. This segment is connected to a 100Base-T switch through an uplink port on the switch which supports 1000Base-T. As an option, one might instead purchase a 100Base-T switch which incorporates a 1000Base-SX port directly (in this case there is no external 1000Base-T segment or media converter). Below the 100Base-T switch are individual pieces of end-user equipment. This network is easily expanded; the center section of the figure shows a possible expansion by adding an additional hub onto one segment from the 100Base-T switch. The hub shown incorporates a bridge to connect 100 MBPS equipment and 10 MBPS equipment.

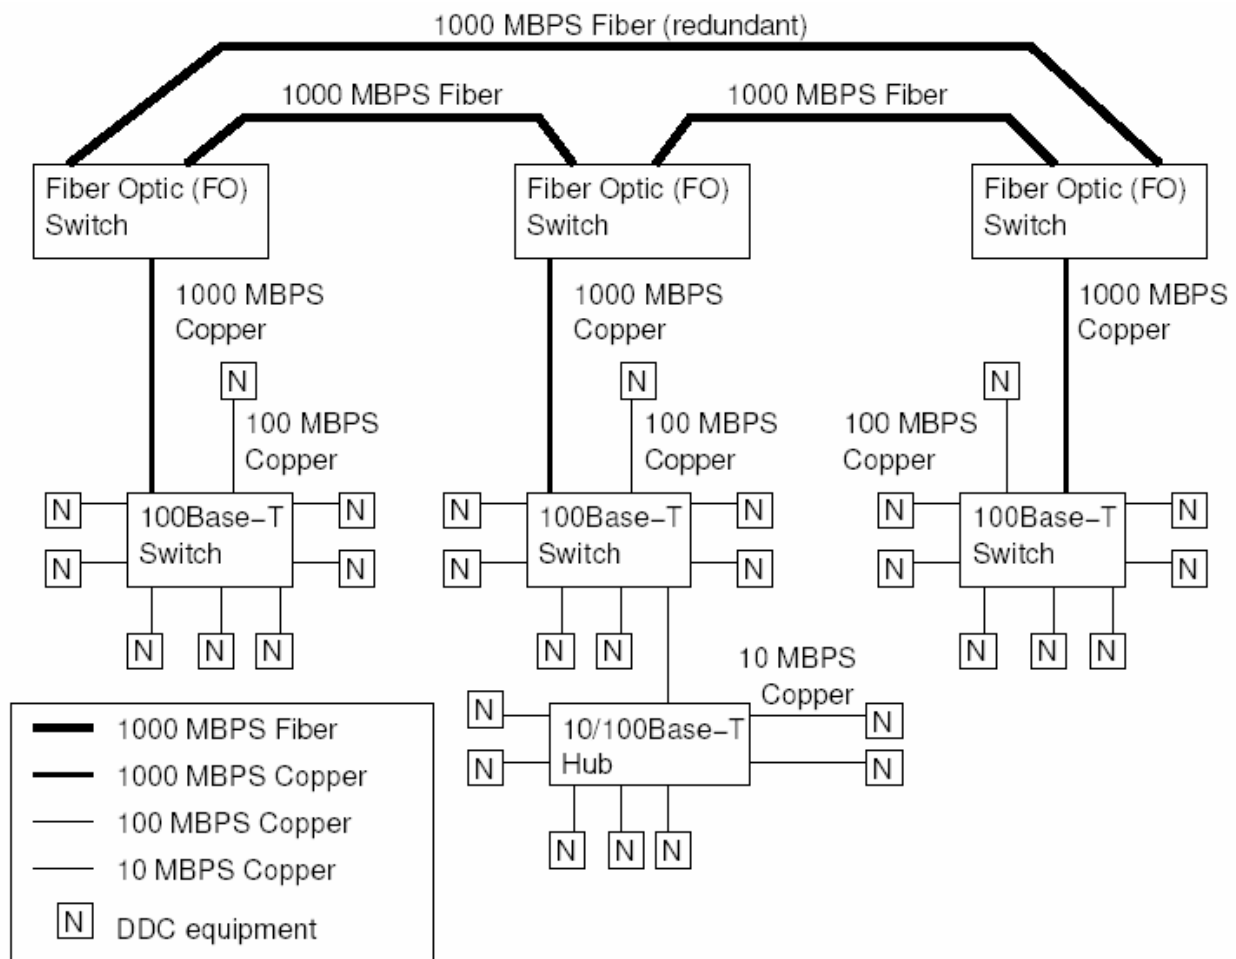


Figure 3-1. Typical Ethernet Network

Figure 3-2 shows a simplified picture of the same network in Figure 3-1, where for the sake of clarity we have drawn the individual networks radiating out from switches as a bus topology. In most cases, the difference between the star and bus topology can be either inferred or ignored.

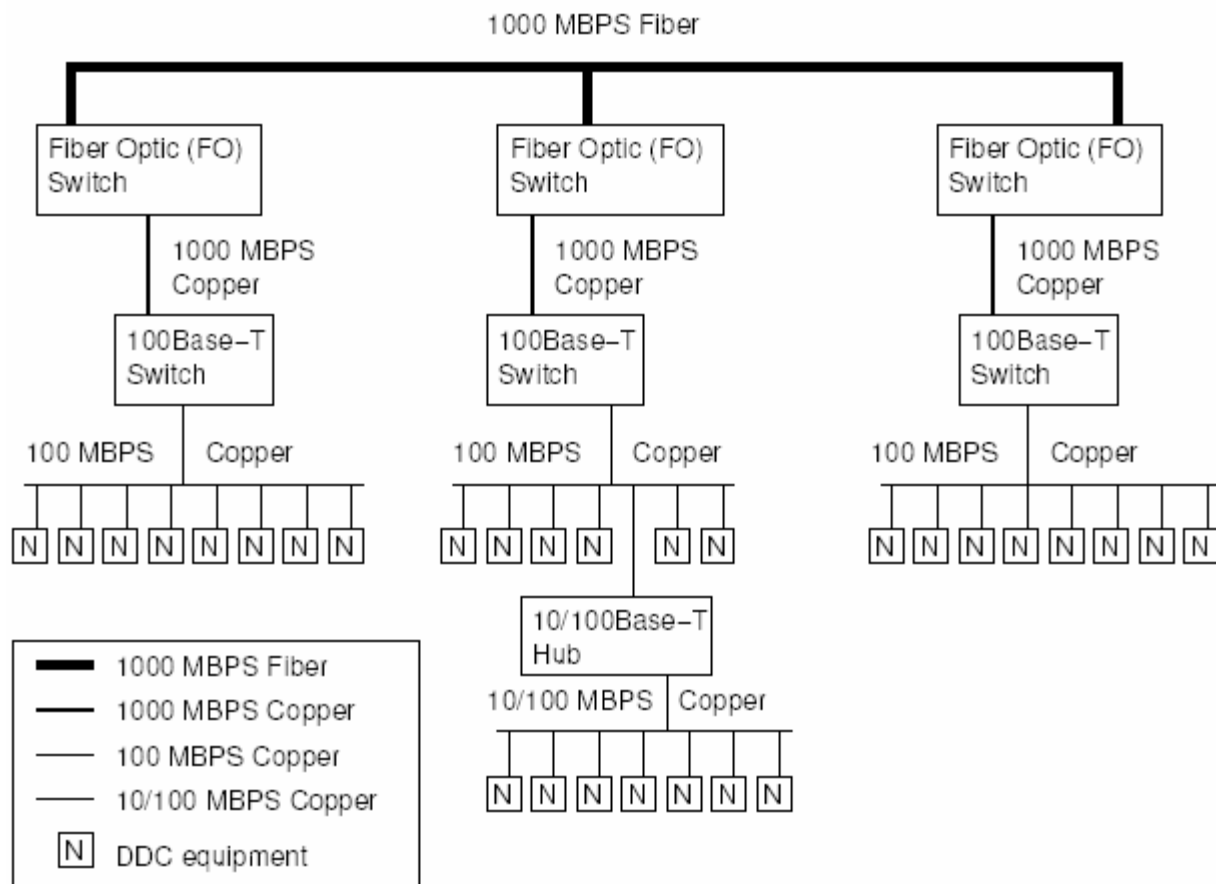


Figure 3-2. Simplified Ethernet Network

An excellent resource for Ethernet is
<http://www.ethermanage.com/ethernet/ethernet.html>

3-2.11 IP Network

IP (Internet Protocol) networks have become the defacto standard for networks ranging from small Local Area Networks (LAN) to the Internet. The IP refers to a very specific protocol, namely the media independent network addressing performed at the Network Layer (layer 3) of the OSI model. This addressing is in the form of a 4-byte number, usually expressed as 4 integers in the range of 0 - 255 (the so-called "dotted-quad" notation). This second layer of addressing provides a number of important benefits over the Ethernet layer 2 addressing (MAC addressing):

- Addresses can be assigned in a logical manner. Machines in some logical grouping can have IP addresses that are close to each other. So, for example, all computers at Camp Swampy might be assigned an IP address in the range of 131.27.0.0 to 131.27.255.255. This makes routing of packets easier: if I have a packet with a destination IP address of 137.27.xxx.yyy (where xxx and yyy are each some number) I would automatically know to route the packet to Camp Swampy, where another machine will route it to it's final destination (based on the value of the xxx and yyy). I don't need to know the ultimate destination of the packet, just that it goes to somewhere in Camp Swampy.

- Addresses are independent of the physical hardware. This permits a global (literally, worldwide) addressing scheme that functions over Ethernet, Asynchronous Transfer Mode (ATM) networks, and other networking hardware. This scheme is even robust enough to permit Virtual Private Networks (VPN) , where the physical media is emulated in software running over a totally independent network. The application software does not need to know what the physical media or hardware address is, it is sufficient to know the IP address.

3-2.12 IP Routers

Devices that route data based on the IP address of the packet are called IP routers, or more simply routers. IP routers are often combined with other functionality in the same hardware. For example, the 100Base-T Switches with 1000Base-T uplink ports shown in the middle of figures 7.1 and 7.2 may also incorporate IP routing.

3-2.13 Other Protocols That Run On IP

Other protocols are commonly associated with IP, but IP is the underlying layer that supports all of the upper layers. Some of these other protocols are TCP (Transmission Control Protocol), UDP (User Datagram Protocol), and ICMP (Internet Control Message Protocol). These protocols function at layers 4 and 5 (depending on the protocol) of the OSI model.

Other protocols such as FTP (File Transfer Protocol), SMTP (Simple Mail Transport Protocol), DNS (Domain Name Service), NTP (Network Time Protocol) , NETBIOS, DHCP (Dynamic Host Configuration Protocol), HTTP (Hypertext Transport Protocol), POP (Post Office Protocol), and NNTP (Network News Transport Protocol) are protocols or applications that run on top of a lower level protocol, usually TCP or UDP.

3-3 IP NETWORK IMPLEMENTATION.

In the case where the UMCS network is required to be completely separate from the IT network, the following technologies should be utilized:

- 100Base-FX fiber optic connections between buildings. Either Single Mode Fiber (SMF) or Multi Mode Fiber (MMF) using duplex SC connectors (possibly MTJR).
- 100Base-T Category 5 copper wiring inside buildings. Due to the low network bandwidth requirements of the UMCS, 1000MBPS Ethernet offers no practical advantages over 100MPBS.
- 100 MBPS Ethernet switch. Switch shall have a minimum of 24 100Base-T ports with RJ-45 connectors and shall be expandable via purchase of additional switches and a cascade connector to a minimum of 96 ports. This expansion shall such that the combined switch is still one switch for network timing and management purposes. Switch shall be installed in appropriate equipment rack

Slower media (10Base-T for example), can be used, but the price difference between it and 100Base-T is minor and, since most manufacturers are moving to 100Base-T, the availability of parts for a 10Base-T network may soon become an issue. Faster media (1000Base-SX or 1000Base-LX) is an option, but the decision cannot be made based on network bandwidth requirements, which can be met by 100 MBPS networking.

Figure 3-3 shows a typical IP network for a UMCS application. Each building is tied to a basewide backbone via a 100Base-FX fiber, media converter, and 100Base-TX

copper Category 5 connection. Inside the building is a Lon-to-IP router and below that is a ANSI-709.1 network with individual LonWorks nodes on it. The center building shows a case where the building was too large for a single network, and the building contains multiple Lon-to-IP routers, connected together with a 100Base-TX network inside the building. At the top, all the buildings are tied together in a single fiber optic switch. The fiber optic switch also has a 100Base-TX port on it, which is connected to a 100Base-TX switch. Finally, personal computers utilized by the UMCS (OWS, LNS server, etc.) are connected to this 100Base-TX switch via Category 5 copper and NICs in the computers.

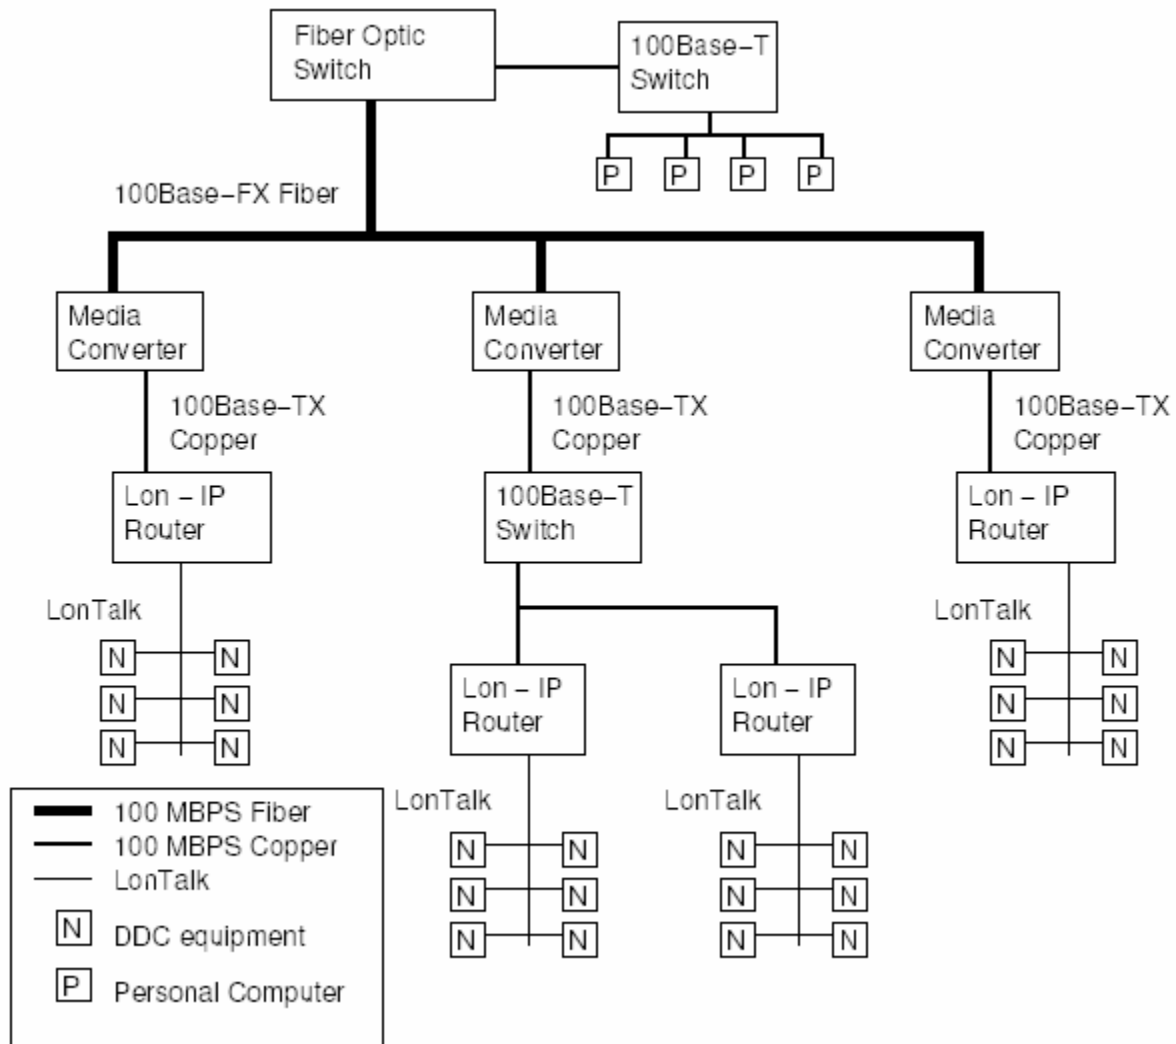


Figure 3-3. Sample UMCS Network

